PCI-EXPRESS SLOT FOR COUPLING PLURAL DEVICES TO A HOST SYSTEM

INVENTOR(S):

PARAG P. MEHTA

PATRICK J. KENNEDY

BACKGROUND

1. Field of the Invention

[0001] The present invention relates to storage systems, and more particularly, to using PCI-Express standard for connecting plural modules.

2. Background of the Invention

[0002] Storage area networks("SAN") are commonly used to store and access data. SAN is a high-speed sub-network of shared storage devices, for example, disks and tape drives. A computer system (may also be referred to as a "host") can access data stored in the SAN.

[0003] Typical SAN architecture makes storage devices available to all servers that are connected using a computer network, for example, a local area network or a wide area network. The term server in this context means any computing system or device coupled to a network that manages network resources. For example, a *file server* is a computer and storage device dedicated to storing files. Any user on the

network can store files on the server. A *print server* is a computer that manages one or more printers, and a *network server* is a computer that manages network traffic. A database *server* is a computer system that processes database queries.

[0004] Various components and standard interfaces are used to move data from host systems to storage devices in a SAN. Fibre channel is one such standard. Fibre channel (incorporated herein by reference in its entirety) is an American National Standard Institute (ANSI) set of standards, which provides a serial transmission protocol for storage and network protocols such as HIPPI, SCSI (small computer system interface), IP, ATM and others. Fibre channel provides an input/output interface to meet the requirements of both channel and network users.

[0005] Host systems often communicate via a host bus adapter ("HBA") using the "PCI" bus interface. PCI stands for Peripheral Component Interconnect, a local bus standard that was developed by Intel Corporation ®. The PCI standard is incorporated herein by reference in its entirety. Most modern computing systems include a PCI bus in addition to a more general expansion bus (e.g. the ISA bus). PCI is a 64-bit bus and can run at clock speeds of 33 or 66 MHz.

[0006] PCI-X is a standard bus that is compatible with existing PCI cards using the PCI bus. PCI-X improves the data transfer rate of PCI from 132 MBps to as much as 1 GBps. The PCI-X standard was developed by IBM®, Hewlett Packard Corporation ® and Compaq Corporation® to increase performance of high bandwidth devices, such as Gigabit Ethernet standard and Fibre Channel Standard, and processors that are part of a cluster.

[0007] PCI-Express (may also be referred to as "PCI-Exp") is another industry standard that is being developed to allow data transfer at 2.5 Gigabits/second and has a layered structure. PCI-Exp provides a dual-simplex channel that is implemented as a transmit and receive pair.

[0008] PCI-Exp link consists of two low-voltage, differentially driven pair of signals, i.e. a transmit pair and a receive pair. A data clock is embedded using an 8b/10b-encoding scheme to achieve high data rates. A PCI-Exp physical layer is used to transport packets between link layers of two PCI-Exp agents. Adding signal pairs to form multiple lanes may linearly scale the bandwidth of a PCI-Exp lane. The current PCI-Exp physical layer can support a x1(single), x2(double), x4(four), x8(eight),x12(twelve), x16(sixteen) and x32(thirty two) lane widths.

[0009] As discussed above, servers to interact with storage sub-systems use adapters. Often multiple adapters are used in complex systems. In order to couple plural adapters to a host system (for example, the host system 101A, Figure 1A), a bridge is required. Figure 1D block diagram shows two adapters (A and B) 106 that are coupled to a bridge 106A allowing host system 101A access to both the adapters A and B. The cost of using bridge 106A is not commercially desirable.

[0010] Therefore, there is a need for a system that allows multiple adapters to be coupled to a host system without using a bridge and preferably using the same PCI-Exp slot.

SUMMARY OF THE INVENTION

[0011] In one aspect of the present invention, a PCI-Express slot for coupling devices to a host system is provided. The slot includes a PCI-Express connector that can couple at least two devices using at least two independent PCI-Express lanes. Four, eight, twelve, sixteen, and/or thirty PCI-Express lanes are used to couple at least two devices, and/or eight PCI-Express lanes are used to couple at least two devices.

[0012] In another aspect of the present invention, a system for coupling plural adapters to a host system

is provided. A motherboard with a PCI-Express slot having a PCI-Express connector that can couple at least two devices using at least two independent PCI-Express lanes.

[0013] In yet another aspect of the present invention, a PCI-Express connector in a PCI-Express slot for coupling devices to a host system is provided, where the connector can couple at least two devices using at least two independent PCI-Express lanes.

[0014] In one aspect of the present invention, a bridge is not needed to couple plural adapters to a host system using the PCI-Exp bus.

[0015] This brief summary has been provided so that the nature of the invention may be understood quickly. A more complete understanding of the invention can be obtained by reference to the following detailed description of the preferred embodiments thereof concerning the attached drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] The foregoing features and other features of the present invention will now be described with reference to the drawings of a preferred embodiment. In the drawings, the same components have the same reference numerals. The illustrated embodiment is

intended to illustrate, but not to limit the invention. The drawings include the following Figures:

[0017] Figure 1A shows a host system with an adapter coupled to a storage subsystem, used according to one aspect of the present invention;

[0018] Figure 1B shows a block diagram with details of an adapter, used according to one aspect of the present invention;

[0019] Figure 1C shows the internal architecture of a host system, used according to one aspect of the present invention;

[0020] Figure 1D shows a top-level block diagram where two adapters are connected to a host system using a bridge;

[0021] Figures 2A and 2B show block diagrams where the same PCI-Exp connector is used to couple more than one adapter, according to one aspect of the present invention; and

[0022] Figure 2C shows a block diagram where the same PCI-Exp connector is used to couple plural adapters, according to one aspect of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0023] To facilitate an understanding of the preferred embodiment, the general architecture and operation of a system using storage devices will be described. The

specific architecture and operation of the preferred embodiment will then be described with reference to the general architecture.

[0024] Figure 1A shows a host system 101A with memory 101 coupled to a SAN 115 that is coupled to storage subsystem 115. In the Figure 1B example, Host 101A has access to storage sub-systems 116 and 118. It is noteworthy that a host system 101A, as referred to herein, may include a computer, server or other similar devices, which may be coupled to storage systems. Host system 101A includes a host processor, random access memory ("RAM"), and read only memory ("ROM"), and other components to communicate with various SAN modules, as described below.

[0025] Figure 1B shows a system 100 that uses a controller /adapter 106 (referred to as "adapter" 106) for communication between a host system 101A with host memory 101 to various storage systems (for example, storage subsystem 116 and 121, tape library 118 and 120) using fibre channel storage area networks 114 and 115.

[0026] Host system 101A communicates with adapter 106 via a PCI-Exp bus 105 through a PCI-Exp interface 107. Adapter 106 includes processors 112 and 109 for the

receive and transmit side, respectively. Processor 109 and 112 may be a RISC processor.

[0027] Transmit path in this context means data coming from host memory 101 to the storage systems via adapter 106. Receive path means data coming from storage subsystem via adapter 106. It is noteworthy, that only one processor can be used for receive and transmit paths, and the present invention is not limited to any particular number/type of processors.

[0028] Adapter 106 also includes fibre channel interface (also referred to as fibre channel protocol manager "FPM") 122 and 113 in receive and transmit paths, respectively. FPM 122 and 113 allow data to move to/from storage systems 116, 118, 120 and 121.

[0029] Adapter 106 includes external memory 108 and 110 and frame buffers 11A and 11B that are used to move information to and from the host to other SAN components.

[0030] Host memory 101 includes a response queue 104 and a request queue 103 to move information to and from host memory 101 using a driver 102.

[0031] Figure 1C is a block diagram showing the internal functional architecture of host system 101A. As shown in Figure 1C, host system 101A includes a microprocessor or central processing unit ("CPU") 124

for executing computer-executable process steps and interfaces with a computer bus 125 (similar to PCI-Exp bus 105). Also shown in Figure 1C is an adapter interface 126 (similar to PCI-Exp interface 107) that interfaces host system 101A with adapter 106. Host system 101A also includes a display device interface 127, a keyboard interface 128, a pointing device interface 132, and a storage device 129 (for example, a disk, CD-ROM or any other device).

[0032] Storage 129 stores operating system program files, application program files, and other files. Some of these files are stored on storage 129 using an installation program. For example, CPU 124 executes computer-executable process steps of an installation program so that CPU 124 can properly execute the application program.

[0033] A random access main memory ("RAM") 130 also interfaces to computer bus 125 to provide CPU 124 with access to memory storage. When executing stored computer-executable process steps from storage 129, CPU 124 stores and executes the process steps out of RAM 130.

[0034] Read only memory ("ROM") 131 is provided to store invariant instruction sequences such as start-up instruction sequences or basic input/output

operating system (BIOS) sequences for operation of a keyboard (not shown).

[0035] Figure 2A shows a block diagram where a motherboard 200 is used in host system 101A and includes a PCI-Exp slot 201. Slot 201 includes a PCI-Exp connector 201A (Figure 2B). A card 203 is placed in slot 201 and connector 201A is used to connect two adapters (A and B) 106 to host system 101A.

[0036] It is noteworthy that although the example in Figures 2A and 2B show two adapters, the present adaptive aspects of the present invention are not limited to any particular number of adapters. For example, in Figure 2C, PCI-Exp slot 201 is shown where two devices 201A and 201B(adapters A and B) are coupled using X4 (i.e. 4 lanes each) in an x4 configuration. Figure 2D also shows PCI-Exp slot 201 that couples four devices 201C-201F using the x2 (two lanes) configuration. Similarly, other devices can be used in x12, x16 and x32 configuration.

[0037] Although the foregoing examples show how adapters in the SAN environment being coupled using the PCI-Exp bus, the present invention is not limited to any particular type of adapter. For example, plural cards in other environments (for example, multi-media,

graphics, or printing) may be coupled using the adaptive aspects of the present invention.

[0038] In one aspect of the present invention, a bridge is not needed to couple plural adapters to a host system using the PCI-Exp bus.

[0039] Although the present invention has been described with reference to specific embodiments, these embodiments are illustrative only and not limiting. Many other applications and embodiments of the present invention will be apparent in light of this disclosure and the following claims.